

THE INVENTION CLAIMED IS:

1. A method for sharing a multiple queue Ethernet adapter comprising:

5 receiving a frame or packet in the adapter;
determining whether the frame or packet is for one or more of a plurality of partitions that share the adapter;
and

10 if the frame or packet is for one or more of the plurality of partitions that share the adapter:

storing the frame or packet in an adapter cache memory;

determining one or more of the plurality of partitions to which the frame or packet is to be sent; and

15 transferring the frame or packet from the adapter cache memory to a receive queue of each of the one or more partitions to which the frame or packet is to be sent.

2. The method of claim 1 further comprising, if the frame
20 or packet is for one or more of the plurality of partitions that share the adapter, generating an interrupt to notify each of the one or more partitions to which the frame or packet is transferred of the frame or packet.

25 3. The method of claim 1 wherein determining one or more of the plurality of partitions to which the frame or packet is to be sent includes:

accessing a table stored in the adapter; and
determining one or more of the plurality of partitions
30 to which the frame or packet is to be sent based on data stored in the table.

4. The method of claim 3 wherein determining one or more of the plurality of partitions to which the frame or packet is to be sent based on data stored in the table includes

5 determining one or more of the plurality of partitions to which the frame or packet is to be sent based on at least one of a MAC address, VLAN ID/MAC address pair and an IP address stored in the table.

10 5. The method of claim 1 wherein determining one or more of the plurality of partitions to which the frame or packet is to be sent includes:

accessing a value stored in a register, the value indicating a type of address to use for determining one or
15 more of the plurality partitions to which the frame or packet is to be sent;

accessing a table stored in the adapter; and

determining one or more of the plurality of partitions to which the frame or packet is to be sent based on the
20 value stored in the register and data stored in the table.

6. A method for sharing a multiple queue Ethernet adapter comprising:

determining whether one or more of a plurality of
25 partitions have a frame or packet to transmit; and

if one or more of the plurality of partitions have a frame or packet to transmit:

selecting a partition from the plurality of partitions that have a frame or packet to transmit;

transferring the frame or packet corresponding to the selected partition from a transmit queue of the selected partition to the adapter cache memory; and transmitting the frame or packet from the adapter.

7. The method of claim 6 wherein determining whether one or more of the plurality of partitions have a frame or packet to transmit includes:

polling a transmit queue corresponding to each of the plurality of partitions; and

determining whether one or more of the plurality of partitions have a frame or packet to transmit based on polling results from one or more of the plurality of partitions.

8. The method of claim 6 wherein selecting a partition from the plurality of partitions that have a frame or packet to transmit includes:

accessing a table stored in the adapter; and selecting a partition from the plurality of partitions that have a frame or packet to transmit based on data stored in the table.

9. The method of claim 8 wherein selecting a partition from the plurality of partitions that have a frame or packet to transmit based on data stored in the table includes selecting a partition from the plurality of partitions that have a frame or packet to transmit based on a priority value stored in the table.

10. A method of sharing a multiple queue Ethernet adapter comprising:

employing a receive queue and a transmit queue for each of a plurality of partitions included in a computer system; and

at least one of transferring a frame or packet from the transmit queue of one of the plurality of partitions to the adapter cache memory, and transferring a frame or packet from the adapter to the receive queue of one of the plurality of partitions.

11. The method of claim 10 wherein transferring a frame or packet from the transmit queue of one of the plurality of partitions to the adapter cache memory includes:

accessing a table stored in the adapter; and transferring a frame or packet from the transmit queue of one of the plurality of partitions to the adapter cache memory based on data stored in the table; and

wherein transferring a frame or packet from the adapter to the receive queue of one of the plurality of partitions includes:

accessing the table stored in the adapter; and transferring a frame or packet from the adapter to the receive queue of one of the plurality of partitions based on data stored in the table.

12. A method of configuring a plurality of partitions of a computer system to share a multiple queue Ethernet adapter comprising:

creating a new partition in the computer system; and

allowing the new partition to share the adapter with one or more other partitions of the computer system.

13. The method of claim 12 wherein allowing the new partition to share the adapter with one or more other partitions of the computer system includes:

selecting the new partition to share the adapter;
sending the address of the selected partition to firmware of the computer system;

employing the firmware to notify a hosting partition that the new partition is allowed to share the adapter; and
updating a table stored in the adapter, the table storing information about the queues corresponding to each partition that shares the adapter.

14. The method of claim 13 further comprising notifying the firmware of an interrupt corresponding to the new partition.

15. The method of claim 13 further comprising updating a table stored in the hosting partition.

16. The method of claim 15 wherein updating the table stored in the adapter includes updating the table stored in the adapter with the table stored in the hosting partition.

17. An apparatus comprising:

a plurality of processors;
a main memory;

a plurality of partitions each including a transmit queue and a receive queue; and

a multiple queue Ethernet adapter coupled to the plurality of processors, main memory and the plurality of partitions, wherein the adapter includes an adapter cache memory, a register, and a table, wherein the plurality of partitions share the adapter and wherein the adapter is adapted to:

receive a frame or packet in the adapter;

determine whether the frame or packet is for one or more of the plurality of partitions that share the adapter; and

if the frame or packet is for one or more of the plurality of partitions that share the adapter:

store the frame or packet in an adapter cache memory;

determine one or more of the plurality of partitions to which the frame or packet is to be sent; and

transfer the frame or packet from the adapter cache memory to the receive queue of each of the one or more partitions to which the frame or packet is to be sent.

18. The apparatus of claim 17 wherein the adapter is further adapted to, if the frame or packet is for one or more of the plurality of partitions that share the adapter, generate an interrupt to notify each of the one or more partitions to which the frame or packet is transferred of the frame or packet.

19. The apparatus of claim 17 wherein the adapter is further adapted to:

access the table stored in the adapter; and

determine one or more of the plurality of partitions to which the frame or packet is to be sent based on data stored in the table.

20. The apparatus of claim 19 wherein the adapter is further adapted to determine one or more of the plurality of partitions to which the frame or packet is to be sent based on at least one of a MAC address, VLAN ID/MAC address pair and an IP address stored in the table.

21. The apparatus of claim 17 wherein the adapter is further adapted to:

access a value stored in the register, the value indicating a type of address to use for determining one or more of the plurality of partitions to which the frame or packet is to be sent;

access the table stored in the adapter; and

determine one or more of the plurality of partitions to which the frame or packet is to be sent based on the value stored in the register and data stored in the table.

22. An apparatus comprising:

a plurality of processors;

a main memory;

a plurality of partitions each including a transmit queue and a receive queue; and

a multiple queue Ethernet adapter, including an adapter cache memory and a table, coupled to the plurality of processors, main memory and the plurality of partitions, and adapted to:

determine whether one or more of the plurality of partitions have a frame or packet to transmit; and

if one or more of the plurality of partitions have a frame or packet to transmit:

5 select a partition from the plurality of partitions that have a frame or packet to transmit;
 transfer the frame or packet corresponding to the selected partition from the transmit queue of the selected partition to the adapter cache memory; and
 10 transmit the frame or packet from the adapter.

23. The apparatus of claim 22 wherein the adapter is further adapted to:

15 poll a transmit queue corresponding to each of the plurality of partitions; and
 determine whether one or more of the plurality of partitions have a frame or packet to transmit based on polling results from one or more of the plurality of
 20 partitions.

24. The apparatus of claim 22 wherein the adapter is further adapted to:

 access the table stored in the adapter; and
 25 select a partition from the plurality of partitions that have a frame or packet to transmit based on data stored in the table.

25. The apparatus of claim 24 wherein the adapter is
 30 further adapted to select a partition from the plurality of

partitions that have a frame or packet to transmit based on a priority value stored in the table.

26. An apparatus comprising:

5 a plurality of processors;
 a main memory;
 a plurality of partitions; and
 a multiple queue Ethernet adapter, including a table,
 coupled to the plurality of processors, main memory and the
 10 plurality of partitions, and adapted to:

 employ a receive queue and a transmit queue for
 each of a plurality of partitions included in a computer
 system; and

 at least one of transfer data from the transmit
 15 queue of one of the plurality of partitions to the adapter
 cache memory, and transfer data from the adapter to the
 receive queue of one of the plurality of partitions.

27. The apparatus of claim 26 wherein the adapter is

20 further adapted to:

 (a) access the table stored in the adapter; and
 transfer data from the transmit queue of one of
 the plurality of partitions to the adapter cache memory
 based on data stored in the table; or

25 (b) access the table stored in the adapter; and
 transfer data from the adapter to the receive
 queue of one of the plurality of partitions based on data
 stored in the table.

28. A computer system for configuring multiple logical partitions to share a multiple queue Ethernet adapter comprising:

a partition management tool;

an apparatus, coupled to the partition management tool, and comprising:

a plurality of processors, wherein the plurality of processors includes a processor for executing firmware;

a main memory;

a plurality of partitions, wherein each partition includes a transmit queue and a receive queue, wherein the plurality of partitions includes a hosting partition and wherein the hosting partition includes a table; and

a multiple queue Ethernet adapter coupled to the plurality of processors, main memory, plurality of partitions and the partition management tool, wherein the adapter includes a table, wherein the plurality of partitions share the adapter, and wherein the hosting partition owns the adapter; and

wherein the computer system is adapted to:

create a new partition in the computer system;

and

allow the new partition to share the adapter with one or more other partitions.

29. The computer system of claim 28 wherein the computer system is further adapted to:

select the new partition to share the adapter;

send the address of the selected partition to firmware of the computer system;

employ the firmware to notify the hosting partition that the new partition is allowed to share the adapter; and

5 update the table stored in the adapter, the table storing information about the transmit and receive queues corresponding to each partition that shares the adapter.

30. The computer system of claim 28 wherein the computer
10 system is further adapted to notify the firmware of an interrupt corresponding to the new partition.

31. The computer system of claim 28 wherein the computer
15 system is further adapted to update the table stored in the hosting partition.

32. The computer system of claim 31 wherein the computer
20 system is further adapted to update the table stored in the adapter with the table stored in the hosting partition.

33. A computer program product comprising:

 a medium readable by a computer, the computer readable medium having computer program code adapted to:

 create a new partition in the computer
25 system; and

 allow the new partition to share the adapter with one or more other partitions of the computer system.

34. The computer program product of claim 33 comprising:

a medium readable by a computer, the computer readable medium having computer program code further adapted to:

select the new partition to share the
5 adapter;

send the address of the selected partition to firmware of the computer system;

employ the firmware to notify a hosting partition that the new partition is allowed to share the
10 adapter; and

update a table stored in the adapter, the table storing information about the queues corresponding to each partition that shares the adapter.

15 35. The method of claim 1 wherein:

determining whether the frame or packet is for one or more of a plurality of partitions that share the adapter includes determining the frame or packet is a broadcast frame or packet; and

20 transferring the frame or packet from the adapter cache memory to a receive queue of each of the one or more partitions to which the frame or packet is to be sent includes transferring the broadcast frame or packet from the adapter cache memory to a receive queue of all of the
25 plurality of partitions that share the adapter.

36. The method of claim 2 wherein generating an interrupt to notify each of the one or more partitions to which the frame or packet is transferred of the frame or packet
30 includes generating a Message Signaling Interrupt (MSI) to

notify each of the one or more partitions to which the frame or packet is transferred of the frame or packet.

37. The method of claim 6 wherein:

5 transferring the frame or packet corresponding to the selected partition from a transmit queue of the selected partition to the adapter cache memory includes determining the frame or packet transferred from the transmit queue of the selected partition to the adapter cache memory is a

10 broadcast frame or packet; and

 transmitting the frame or packet from the adapter includes transferring the broadcast frame or packet to the receive queue of all partitions that share the adapter except for the selected partition.

15 38. The apparatus of claim 17 wherein the adapter is further adapted to:

 determine the frame or packet is a broadcast frame or packet; and

20 if the frame or packet is for one or more of the plurality of partitions that share the adapter, transfer the broadcast frame or packet from the adapter cache memory to a receive queue of all of the plurality of partitions that share the adapter.

25 39. The apparatus of claim 18 wherein the adapter is further adapted to, if the frame or packet is for one or more of the plurality of partitions that share the adapter, generate a Message Signaling Interrupt (MSI) to notify each
30 of the one or more partitions to which the frame or packet is transferred of the frame or packet.

40. The apparatus of claim 22 wherein the adapter is further adapted to:

if one or more of the plurality of partitions have a frame or packet to transmit:

determine the frame or packet transferred from the transmit queue of the selected partition to the adapter cache memory is a broadcast frame or packet; and

transfer the broadcast frame or packet to the receive queue of all partitions that share the adapter except for the selected partition.

41. The method of claim 6 wherein transmitting the frame or packet from the adapter includes transmitting the frame or packet using a network connection or transmitting the frame or packet to one or more of the plurality of partitions.

42. The apparatus of claim 22 wherein the adapter is further adapted to transmit the frame or packet using a network connection or transmit the frame or packet to one or more of the plurality of partitions.

43. A method for sharing a multiple queue Ethernet-protocol adapter within a multiple processor environment having multiple partitions comprising:

receiving a frame or packet in the adapter; and
determining whether the frame or packet is for one or more of the plurality of partitions that share the adapter.

44. The method of claim 43 wherein determining whether the frame or packet is for one or more of the plurality of partitions that share the adapter includes:

accessing an address included in the received frame or
5 packet; and

comparing the accessed address to one or more addresses assigned to the adapter.